

Propositional Abduction with Implicit Hitting Sets

Alexey Ignatiev, Antonio Morgado, and Joao Marques-Silva

LaSIGE, Faculty of Science, University of Lisbon, Portugal, email:
 {aignatiev,ajmorgado,jpms}@ciencias.ulisboa.pt

Abstract. Logic-based abduction finds important applications in artificial intelligence and related areas. One application example is in finding explanations for observed phenomena. Propositional abduction is a restriction of abduction to the propositional domain, and complexity-wise is in the second level of the polynomial hierarchy. Recent work has shown that exploiting implicit hitting sets and propositional satisfiability (SAT) solvers provides an efficient approach for propositional abduction. This paper investigates this earlier work and proposes a number of algorithmic improvements. These improvements are shown to yield exponential reductions in the number of SAT solver calls. More importantly, the experimental results show significant performance improvements compared to the the best approaches for propositional abduction.

1 Introduction

Logic-based abduction finds relevant applications in artificial intelligence and related areas [9, 12, 18–20, 25–27, 33, 49–52, 54, 58–60]. Given a background theory and a set of manifestations and a set of hypotheses, abduction seeks to identify a cost-minimum set of hypotheses which explain the manifestation and are consistent given the background theory. Propositional abduction is hard for the second level of the polynomial hierarchy, but finds a growing number of applications [58, 60]. Noticeable examples of where propositional abduction algorithms can be applied include abductive inference [15–17, 55], logic programming [34, 36, 41], knowledge bases updates [35, 59], security protocols verification [1], and constraint optimization [21, 22], among many others.

Given the complexity class of propositional abduction, it is conceptually simple to solve the problem with a linear (or logarithmic) number of calls to a Σ_2^P -oracle, e.g. a oracle for quantified Boolean formulae (QBF) with one quantifier alternation. Unfortunately, in practice QBF solvers are not as efficient as SAT solvers, and do not scale as well. As a result, recent work [58] on solving propositional abduction focused on using calls to SAT oracles instead of QBF oracles, following a trend also observed for solving QBF [28, 29, 31, 32]. This recent work on solving propositional abduction is motivated by the practical success of implicit hitting set algorithms in a number of different settings [4, 10, 11, 13, 23, 24, 28, 29, 31, 32, 37–39, 45, 53, 58, 61].

The contributions of this paper can be summarized as follows. The paper revisits QBF models for solving propositional abduction and proposes a Quantified MaxSAT (QMaxSAT) [23, 24] model for propositional abduction. Moreover, the paper notes that the MaxHS [13] approach for MaxSAT can be readily applied to QMaxSAT by replacing the oracle used. The paper then investigates the application of implicit hitting sets to solving propositional abduction [58] and identifies a number of algorithmic improvements. This leads to a new algorithm, Hyper, for solving propositional abduction. This new algorithm is shown to significantly outperform the current state of the art, solving a large number of instances that could not be solved with existing solutions. More importantly, the paper shows that the algorithmic improvements proposed can save an exponential number of iterations when compared with the current state of the art [58].

The paper is organized as follows. Section 2 introduces the definitions used throughout the paper, and also overviews related work. Section 3 revisits a QBF model for abduction, which is then used for developing a number of alternative approaches for solving propositional abduction. Among these, the paper proposes improvements to recent work [58], which are shown to yield exponential reductions on the number of SAT oracle calls. Section 4 analyzes the experimental results, running existing and the proposed algorithms on existing problem instances [58]. Section 4 also provides experimental evidence that the proposed algorithms for propositional abduction can save an exponential number of oracle calls in different settings. Section 5 concludes the paper, and identifies possible research directions.

2 Preliminaries

This section introduces the notation and definitions used throughout the paper.

2.1 Satisfiability

Standard propositional logic definitions apply (e.g. [7]). CNF formulas are defined over a set of propositional variables. A CNF formula (or theory) T is a propositional formula represented as a conjunction of clauses, also interpreted as a set of clauses. A clause is a disjunction of literals, also interpreted as a set of literals. A literal is a variable or its complement. The set of variables of a theory T is denoted $X \triangleq \text{var}(T)$. The dependency of T on X can be made explicit by writing $T(X)$. Where convenient, a formula can be rewritten with a fresh set of variables, e.g. we can replace X by Y , writing $T(Y)$. Conflict-driven clause learning (CDCL) SAT solvers are summarized in [7]. Throughout the paper, SAT solvers are viewed as oracles. Given a CNF formula F , a SAT oracle decides whether F is satisfiable, and returns a satisfying assignment μ if F is satisfiable. A SAT oracle can also return a subset of the clauses (i.e. an unsatisfiable core $U \subseteq F$) if F is unsatisfiable.

CNF formulas are often used to model overconstrained problems. In general, clauses in a CNF formula are characterized as hard, meaning that these must

be satisfied, or soft, meaning that these are to be satisfied, if at all possible. A weight can be associated with each soft clause, and the goal of maximum satisfiability (MaxSAT) is to find an assignment to the propositional variables such that the hard clauses are satisfied, and the sum of the satisfied soft clauses is maximized. Branch-and-bound algorithms for MaxSAT are overviewed in [7]. Recent work on MaxSAT investigated core-guided algorithms [3, 46] and also the use of implicit hitting sets [13].

In the analysis of unsatisfiable CNF formulas, a number of definitions are used. Given an unsatisfiable CNF formula F , a minimal unsatisfiable subset (MUS) $M \subseteq F$ is both unsatisfiable and irreducible. Given an unsatisfiable CNF formula F , a minimal correction subset (MCS) $C \subseteq F$ is both irreducible and its complement is satisfiable. Given an unsatisfiable CNF formula F , a maximal satisfiable subset (MSS) S is the complement of some MCS of F . A largest MSS is a solution to the MaxSAT problem.

Additionally, it is well-known [40, 56] that MUSes and MCSes are connected by a *hitting set duality*. Given a collection Γ of sets from a universe \mathbb{U} , a hitting set h for Γ is a set such that $\forall S \in \Gamma, h \cap S \neq \emptyset$. A hitting set h is *minimal* if none of its subset is a hitting set. It is straightforward to extend the previous definitions to the case where there are hard clauses.

Quantified Boolean formulas (QBFs) are an extension of propositional logic with *existential* and *universal* quantifiers (\forall, \exists) [7]. A QBF can be in *prenex closed* form $Q_1x_1 \dots Q_nx_n. \varphi$, where $Q_i \in \{\forall, \exists\}$, x_i are distinct Boolean variables, and φ is a Boolean formula over the variables x_i and the constants 0 (false), 1 (true). The sequence of quantifiers in a QBF is called the *prefix* and the Boolean formula the *matrix*. The semantics of QBF is defined recursively. A QBF $\exists x_1 Q_2x_2 \dots Q_nx_n. \varphi$ is true iff $Q_2x_2 \dots Q_nx_n. \varphi|_{x_1=1}$ or $Q_2x_2 \dots Q_nx_n. \varphi|_{x_1=0}$ is true. A QBF $\forall x_1 Q_2x_2 \dots Q_nx_n. \varphi$ is true iff both $Q_2x_2 \dots Q_nx_n. \varphi|_{x_1=1}$ and $Q_2x_2 \dots Q_nx_n. \varphi|_{x_1=0}$ are true. To decide whether a given QBF is true or not, is known to be PSPACE-complete [7].

2.2 Propositional Abduction

A propositional abduction problem (PAP) is a 5-tuple $P = (V, H, M, T, c)$. V is a finite set of variables. H , M and T are CNF formula representing, respectively, the set of hypotheses, the set of manifestations, and the background theory. c is a cost function associating a cost with each clause of H , $c : H \rightarrow \mathbb{R}^+$.

Given a background theory T , a set $S \subseteq H$ of hypotheses is an explanation (for the manifestations) if: (i) S entails the manifestations M (given T); and (ii) S is consistent (given T). The propositional abduction problem consists in computing a minimum size explanation for the manifestations subject to the background theory.

Definition 1 (Explanations for P [58]). Let $P = (V, H, M, T, c)$ be a PAP. The set of explanations of P is given by the set $\text{Expl}(P) = \{S \subseteq H \mid T \wedge S \not\models \perp, T \wedge S \models M\}$. The minimum-cost solutions of P are given by $\text{Expl}_c(P) = \arg\min_{E \in \text{Expl}(P)} (c(E))$.

Input: F WCNF formula
Output: $(\mu, \text{Cost}(\mu))$ MaxSAT assignment and cost

```

1 begin
2    $K \leftarrow \emptyset$ 
3   while true do
4      $h \leftarrow \text{MinimumHS}(K)$ 
5      $(st, \mu) \leftarrow \text{SAT}(F \setminus h)$ 
      // If  $st$ , then  $\mu$  is an assignment
      // Otherwise,  $\mu$  is a core
6     if  $st$  then return  $(\mu, \text{Cost}(\mu))$   $K \leftarrow K \cup \{\mu\}$ 
7 end

```

Algorithm 1: The MaxHS algorithm [13]

The complexity of logic-based abduction has been investigated in a number of works [9, 18], and is surveyed in [58]. Checking whether $S \subseteq H$ is an explanation for a PAP is D^{P} -complete. Deciding the existence of some explanation is Σ_2^{P} -complete. Finding a minimum-size explanation can be achieved with a linear number of calls to a Σ_2^{P} oracle or, if the costs are polynomially bounded, with a logarithmic number of calls to a Σ_2^{P} oracle.

Example 1 (Example abduction instance.). Consider the propositional abduction problem instance $P = (V, H, M, T, c)$ with the set of variables V , the set of hypotheses H , the manifestations M , and the background theory T given by,

$$\begin{aligned}
V &= \{x_1, x_2, x_3, x_4\} \\
H &= \{(x_1), (x_2), (x_3)\} \\
M &= \{(x_4)\} \\
T &= \{(\neg x_1 \vee x_4), (\neg x_2 \vee \neg x_3 \vee x_4)\}
\end{aligned} \tag{1}$$

The (propositional) abduction problem for this example is to find a minimum cost subset S of H , such that (i) S is consistent with T (i.e. $T \wedge S \not\models \perp$); and (ii) S and T entail M (i.e. $T \wedge S \models M$). For this instance of propositional abduction, the minimum cost explanation is then $S = \{(x_1)\}$.

2.3 Related Work

This paper builds on recent work on algorithms for propositional abduction [58], which builds on earlier work on solving maximum satisfiability with implicit hitting set algorithms [13]. This work is also tightly related with the body of work on handling hitting sets implicitly [10, 11, 13, 37, 45, 58], which is also tightly related with implicit hitting set dualization [4, 38, 39, 53, 61], but also with abstraction refinement in QBF solving and optimization [23, 24, 28, 29, 31, 32].

The use of implicit hitting sets for solving MaxSAT is embodied by MaxHS (e.g. see [13]), which is summarized in Algorithm 1. The algorithm computes minimum hitting sets of a set of sets, each of which represents an unsatisfiable

Input: PAP $P = (V, H, M, T, c)$
Output: Minimum cost explanation S

```

1 begin
2    $K \leftarrow \emptyset$ 
3    $S \leftarrow \emptyset$ 
4    $(st, \mu) \leftarrow \text{SAT}(T \wedge H \wedge (\neg M))$ 
5   if  $st$  then return  $\emptyset$  while  $S \neq H$  do
6      $(st, \mu) \leftarrow \text{SAT}(T \wedge S \wedge (\neg M))$ 
7     if  $st$  then  $K \leftarrow K \cup \{h \in H \mid \mu(h) = 0\}$  else
8        $(st, \mu) \leftarrow \text{SAT}(T \wedge S)$ 
9       if not  $st$  then  $K \leftarrow K \cup \{H \setminus S\}$  else return  $S$ 
10     $S \leftarrow \text{MinimumHS}(K)$ 
11  return  $\emptyset$ 
12 end

```

Algorithm 2: The AbHS/AbHS+ abduction algorithm [58]

subformula of the target formula. This essentially exploits Reiter's [56] well-known hitting set relationship between MCSes and MUSes [6, 8, 40], where an MCS is a minimal hitting set of the MUSes and vice-versa. Moreover, since a minimum hitting set is being computed, we are in search of the smallest MCS, i.e. the MaxSAT solution. In case there are hard clauses, MaxHS needs to take this into consideration, including checking the consistency of the hard clauses. Besides MaxHS, recent work on MaxSAT solving is based on iterative unsatisfiable core identification [3, 46].

Recent work on propositional abduction builds on MaxHS and proposes a novel algorithm, AbHS/AbHS+. AbHS mimics MaxHS in that the algorithm iteratively computes minimum cost hitting sets, which identify a subset $S \subseteq H$. This set S is then used for checking whether it represents an explanation of the propositional abduction problem. Since it is a minimum hitting set then, if the conditions hold, it is a minimum-cost explanation. AbHS is summarized in Algorithm 2. As in MaxHS, the algorithm iteratively computes minimum hitting sets using an ILP solver (line 10). The outcome is a subset S of H . The abduction conditions are checked with two distinct SAT oracle calls. One oracle call checks whether $T \wedge S \wedge (\neg M)$ is inconsistent, i.e. whether $T \wedge S \models M$. If the formula is satisfiable, then the set of sets to hit (K) is updated with another set, of the clauses in H falsified by the computed satisfying assignment. If $T \wedge S \wedge (\neg M)$ is inconsistent, then a second oracle call checks whether $T \wedge S$ is consistent. If it is, then a minimum-cost explanation has been identified. Otherwise, AbHS creates a hitting set by requiring that some non-selected clause of H be selected in subsequently computed minimum hitting sets. For the AbHS+ variant [58], the set added can be viewed as the complements of the literals selecting each clause in S , i.e. at least some clause in S must not be picked.

There is a vast body of work on exploiting implicit hitting sets. The concept of exploiting implicit hitting sets is intended to mean that, instead of starting from an explicit representation of the complete set of hitting sets, hitting sets are computed on demand, as deemed necessary by the problem being solved. An earlier example of exploiting implicit hitting sets is the work of Bailey and Stuckey [6], in the concrete application to hitting set dualization. The concept was re-introduced more recently [10, 13, 37], and then applied in a number of different settings. Among this vast body of work, as will become clear throughout the paper, our work can be related with abstraction refinement ideas used in recent expansion-based QBF solvers, namely RAReQS [24, 28, 31] and quantified optimization extensions [23, 29], but now in the context of handling implicit hitting sets.

3 Algorithms for Propositional Abduction

This section overviews different algorithms for solving propositional abduction, all of which are based on reducing the problem to QBF.

3.1 QBF Model for Abduction

Given a PAP $P = (V, H, M, T, c)$, the problem of deciding whether some set S is an explanation can be reduced to QBF. $S \subseteq H$ is an explanation of P iff:

$$\exists_X T(X) \wedge S(X) \wedge \forall_Y \neg(T(Y) \wedge S(Y) \wedge \neg M(Y)) \quad (2)$$

is true. (Observe X and Y denote sets of variables, thus highlighting that different sets of variables are used.) (2) can be rewritten as follows:

$$\exists_X \phi(X) \wedge \forall_Y \psi(Y), \quad (3)$$

where $\phi = T \wedge S$ and $\psi = \neg(T \wedge S \wedge \neg M)$.

As indicated in Section 2, the goal of propositional abduction is to find a minimum cost explanation, i.e. to pick a minimum cost set $S \subseteq H$ that is an explanation of P .

The problem of finding a minimum cost explanation of P can be reduced to quantified maximum satisfiability [24] (QMaxSAT). Associate a variable r_i with each clause $C_i \in H$, and create a set H' where each clause $C_i \in H$ is replaced by $(r_i \vee C_i)$, to enable relaxing the clause. Let R denote the set of the r_i (relaxation) variables, with $|R| = |H|$. H' serves to create a modified QBF:

$$\exists_R \exists_X T(X) \wedge H'(R, X) \wedge \forall_Y \neg(T(Y) \wedge H'(R, Y) \wedge \neg M(Y)) \quad (4)$$

As before, (4) can be rewritten as follows:

$$\exists_R \exists_X \phi(X, R) \wedge \forall_Y \psi(Y, R) \quad (5)$$

The above QBF can be transformed into prenex normal form, and represents the hard part of the QMaxSAT problem. Moreover, the fact that the goal is to

compute a minimum cost explanation of P is modeled by adding a soft clause $(\neg r_i)$, with cost $c(C_i)$, for each $r_i \in R$. Each soft clause denotes a preference not to include the associated clause in H in the computed explanation.

Example 2. With respect to the PAP from example [Example 1](#), the QBF associated with the hard part of the QMaxSAT problem is:

$$\begin{aligned} & \exists_{r_1, r_2, r_3} \exists_{x_1, x_2, x_3, x_4} \\ & (\neg x_1 \vee x_4) \wedge (\neg x_2 \vee \neg x_3 \vee x_4) \wedge \\ & (r_1 \vee x_1) \wedge (r_2 \vee x_2) \wedge (r_3 \vee x_3) \\ & \forall_{y_1, y_2, y_3, y_4} \\ & \neg[(\neg y_1 \vee y_4) \wedge (\neg y_2 \vee \neg y_3 \vee y_4) \wedge \\ & (r_1 \vee y_1) \wedge (r_2 \vee y_2) \wedge (r_3 \vee y_3) \wedge (\neg y_4)] \end{aligned} \quad (6)$$

with the soft clauses being $\{(\neg r_1), (\neg r_2), (\neg r_3)\}$.

The QMaxSAT formulation can be used to develop a number of alternative approaches for solving PAP. These approaches are detailed in the next sections.

Observe that, if the propositional abduction problem is not trivial to solve, then the QBF (2) is false for $S = \emptyset$ and $S = H$.

3.2 Abduction with QMaxSAT

Similarly to MaxSAT, a number of algorithms can be envisioned for solving QMaxSAT. These are analyzed in the subsections below, taking into account the specific structure of the reduction of propositional abduction to QMaxSAT.

Iterative QBF Solving A standard approach for solving MaxSAT is iterative SAT solving [46]. Similarly, we can use iterative QBF solving for QMaxSAT. At each step, and given cost k , the following pseudo-Boolean constraint is used:

$$\text{PB}(R, k) \triangleq \left(\sum_{C_i \in H} c(C_i) r_i \leq k \right) \quad (7)$$

For some positive k , the QBF (5) can be used for iteratively QBF solving as follows:

$$\exists_R \text{PB}(R, k) \wedge \exists_X \phi(X, R) \wedge \forall_Y \psi(R, Y) \quad (8)$$

Clearly, binary search can be used to ensure a linear (or logarithmic, depending on whether the costs are bounded) number of Σ_2^P oracle calls [18]. In practice, most QBF solvers expect clausal representations. Clausification introduces one additional level of quantification. Typically, each quantification level makes a QBF formula harder to decide. And thus in practice, QBF solvers scale worse than SAT solvers, and so this approach is unlikely to scale for large propositional abduction problem instances.

Core-Guided QBF Solving Core-guided algorithms [3, 46] represent another approach for solving QMaxSAT. Many variants of core-guided MaxSAT algorithms have been proposed in recent years [3, 46].

Given the reduction of propositional abduction to QMaxSAT, any core-guided MaxSAT algorithm can be used, provided a core-producing QBF solver is used [24, 42].

Nevertheless, for the abduction problem the use of alternative MaxSAT solving approaches, based on MaxHS [13] is amenable to efficient optimizations, which solely use SAT solvers [58].

Exploiting MaxHS A recent approach for MaxSAT solving is MaxHS [13], which exploits integer linear programming (ILP) solving, resulting in simpler SAT oracle calls, at the cost of a possibly exponentially larger number of oracle calls. Recall that the MaxHS approach for solving MaxSAT is outlined in Algorithm 1.

A straightforward solution for solving QMaxSAT is to replace the SAT oracle by a QBF oracle in MaxHS. This approach is referred to as QMaxHS, and it was implemented on top of DepQBF, the known QBF solver which is capable of reporting unsatisfiable cores if the input QBF is false [42]. It should be noted (and it is also mentioned in Section 4) that the implementation of QMaxHS performs quite bad (it cannot solve any benchmark instances considered in Section 4). A possible explanation of this is that the QBF formulas, which are iteratively solved by the QBF solver, are too hard even though the original idea of MaxHS-like algorithms is to get (many) simple calls to the oracle. This suggests that implementing core-guided QMaxSAT algorithms would not pay off as well since the QBF formulas in core-guided QMaxSAT are much harder to deal with. Recent work on propositional abduction [58] proposed a MaxHS-like approach, but the QBF oracle call was replaced by two SAT solver calls, which is expected to outperform QMaxHS.

3.3 Exploiting Implicit Hitting Sets

The use of implicit hitting sets for abduction was proposed in recent work [58]. This work can be viewed as extending the MaxHS algorithm for MaxSAT [13], which is based on implicit enumeration of hitting sets. In contrast to MaxHS, instead of one SAT oracle call, AbHS [58] uses two SAT oracle calls, one to check entailment of M by $T \wedge S$ and another to check the consistency of $T \wedge S$. A variant of AbHS, AbHS+, differs on which sets are added to the hitting set representation. Whereas the connection of MaxHS and AbHS with implicit hitting sets is clear, the approach used in AbHS+ can be viewed as adding both only positive clauses and only negative clauses to hit, and so the connection with hitting sets is less evident. An alternative way of explaining AbHS/AbHS+ is to consider (5). The ILP solver is used for computing some minimum cost hitting set, which represents a set of clauses $S \subseteq H$. Then, one SAT oracle call checks $\exists_X \phi(X)$, given S , and another SAT oracle call checks $\forall_Y \psi(Y)$, also given S . (Observe that this second formula corresponds to checking unsatisfiability.)

This explanation of how AbHS/AbHS+ works is investigated in greater detail below.

In the following an alternative approach for propositional abduction is developed which, similarly to AbHS/AbHS+, is also based on handling implicit hitting sets, but which is shown to yield exponential reductions on the number of oracle calls in the worst case. The new algorithm, Hyper, shares similarities with MaxHS and also with AbHS/AbHS+ in that minimum hitting sets are also computed, and an implicit representation of the hitting sets is maintained. However, and in contrast with AbHS/AbHS+, Hyper analyzes the structure of the problem formulation, and develops a number of optimizations that exploit that formulation.

The propositional abduction problem formulation can be presented in a slightly modified form, i.e. to find a smallest cost set $S \subseteq H$, consistent with T (i.e. a T -consistent set S), which, together with T , entails M . Consider a T -consistent candidate set $S \subseteq H$. If $T \wedge S \not\models M$, then the formula $T \wedge S \wedge (\neg M)$ is satisfiable and the satisfying assignment returned by the SAT oracle is a *counterexample* explaining why the selected set S is such that $T \wedge S \not\models M$. Moreover, this satisfying assignment can be used for revealing a (possibly subset-minimal) set of clauses in $H \setminus S$ which are falsified. Clearly, one of these falsified clauses must be included (i.e. *hit*) in any T -consistent set $S \subseteq H$ which, together with T , will entail M . Thus, from each T -consistent candidate set $S \subseteq H$ which, together with T , does not entail M , we can identify a set of clauses, from which at least one must be picked, in order to pick another T -consistent set $S \subseteq H$, such that eventually $T \wedge S \models M$.

The approach outlined in the paragraph above, although apparently similar to the description of AbHS/AbHS+, reveals a number of significant insights. First, checking the consistency of S with T can be carried out *concurrently* with the selection of S itself. This is also apparent from the QBF formulation (2), in that all existential quantifiers can be aggregated and handled simultaneously. Thus, each minimum hitting set S is computed while guaranteeing that $T \wedge S$ holds. More importantly, after each set S is picked, it is only necessary to check whether $T \wedge S \models M$, and this can be done with a *single* SAT oracle call.

Concretely, the next minimum cost hitting set, given the already identified sets to hit, is computed guaranteeing that the existential part of (4) is satisfied:

$$\exists_R \exists_X T(X) \wedge H'(R, X) \quad (9)$$

The selected set S , identified by the assignment to the R variables, is then used for checking the satisfiability of the second component of QBF (2):

$$\forall_Y \neg (T(Y) \wedge S(Y) \wedge \neg M(Y)) \quad (10)$$

Observe that this can be decided with a SAT oracle call.

Thus, a careful analysis of the problem formulation enables improving upon AbHS and AbHS+, specifically by eliminating one SAT oracle call per iteration. However, as shown in the next section, the new algorithm can save an exponentially large number of SAT oracle calls when compared with AbHS/AbHS+.

Input: PAP $P = (V, H, M, T, c)$
Output: Minimum cost explanation S

```

1 begin
2    $(H', R) \leftarrow \text{RelaxCls}(H)$ 
3    $B \leftarrow T \wedge M \wedge H'$ 
4    $A \leftarrow \emptyset$ 
5   while true do
6      $(st, h) \leftarrow \text{MinimumHS}(A, B)$ 
7     if not  $st$  then return  $\emptyset$   $S \leftarrow \{C_i \in H' \mid r_i \in h\}$ 
8      $(st, \mu) \leftarrow \text{SAT}(T \wedge S \wedge (\neg M))$ 
9     if not  $st$  then return  $S$   $W \leftarrow \text{PickFalseCls}(H \setminus S, \mu)$ 
10     $Y \leftarrow \text{GetRelaxationVars}(W)$ 
11     $A \leftarrow A \cup Y$ 
12 end

```

Algorithm 3: Organization of Hyper

Besides the aggregation of the existential quantifiers, additional optimizations can be envisioned. Propositional abduction seeks a minimum cost set $S \subseteq H$ such that $T \wedge S \models M$. Ideally, one would prefer not to select a set $S \subseteq H$ such that $T \wedge S \not\models M$. Observe that $T \wedge S \models M$ implies that $T \wedge S \wedge M$ holds, but the converse is in general not true. Thus, M can be added to (9), resulting in requiring that $T \wedge S \wedge M$ be consistent when selecting the set S . The inclusion of M when picking a minimum hitting set can also reduce the number of oracle calls exponentially. This is also investigated in the next section.

The new Hyper algorithm for propositional abduction is shown as [Algorithm 3](#). Clauses in H are relaxed, to allow each clause $C_i \in H$ to be picked when the associated relaxation variable r_i is assigned value 1. The minimum hitting sets are computed for the set of sets A , subject to a background theory B , which conjoins T , M and the relaxed clauses of H . In Hyper minimum hitting sets are computed with a MaxSAT solver [47], since the hard part (containing T , M and H') plays a significant role in deciding consistency. If all hitting sets have been (implicitly) tried unsuccessfully, then the algorithm terminates and returns \emptyset . If not, and if $T \wedge S \models M$, then the algorithm terminates and returns the computed set S . Otherwise, a subset of the clauses in $H \setminus S$, which are falsified by the computed satisfying assignment μ , is identified, and the associated relaxation variables are used to create another set to hit, i.e. one of those clauses must be included in any selected set S .

Additional Optimizations A few additional optimizations are possible, which can be expected to have some impact in the performance of Hyper. These are discussed next. The first two optimizations are implemented in a variant of Hyper, Hyper*. The other optimizations are analyzed to explain why performance improvements are not expected to be significant.

One optimization is to perform *partial reduction* of the counterexamples computed in lines 11–12 of [Algorithm 3](#). Recall that counterexamples, i.e. sets that need to be hit next time, comprise clauses of $H \setminus S$ that are falsified by a model μ of $T \wedge S \wedge (\neg M)$. Thus, the counterexamples can be seen as correction subsets for the partial CNF formula $T \wedge H \wedge (\neg M)$, where $T \wedge (\neg M)$ is the hard part and H is the soft part. Observe that instead of computing *any* correction subset, one may want to reduce it to get a subset-minimal correction subset (an MCS), i.e. to try to minimize the number of falsified clauses in $H \setminus S$. In Hyper* this is done using the standard linear search algorithm [44], which iterates through the falsified clauses and tries to satisfy them. In order not to spend too much time on doing the reduction, the version of Hyper* presented here iterates only over $0.2 \times m$ falsified clauses of the initial counterexample starting from the clauses of the smallest weight, where m is the size of the initial counterexample.

A second optimization is to start by computing a fixed number of minimum hitting sets. Given that $T \wedge H \wedge (\neg M)$ is inconsistent, one can enumerate MCSes of this formula, which must be hit, so that one can eventually prove that there exists some set S with $T \wedge S \models M$. In Hyper, 100 MCSes of $T \wedge H \wedge (\neg M)$ are computed before starting the process of generating candidate sets S . These MCSes are computed by size, using MaxSAT-based MCS enumeration [40, 48].

A third optimization respects the clauses in M . Any $C_j \in M$, such that $T \models C_j$, can be removed from M . In a preprocessing step, each clause in M is checked for entailment with respect to T . Any entailed clause is removed. This technique reduces the practical hardness of the formulas checked for unsatisfiability. Since most of the running time of Hyper is spent on computing minimum hitting sets, the impact of the technique is expected to be marginal.

A fourth, and final optimization respects the clauses in H . Any $C_j \in H$ that $T \models C_j$, can also be removed from H , as it will not be included in any minimum cost hitting set. It should be noted that the gains of this technique should be also marginal. Since by construction $T \wedge S$ is consistent, and the only computed counterexamples satisfy $T \wedge S \wedge (\neg M)$, then any clause $C_j \in H$ with $T \models C_j$ will also be satisfied. Since, the counterexamples only consider falsified clauses, then any clause $C_j \in H$ entailed by T will *never* be included in a set to be hit.

Exponential Reductions in Oracle Calls This section argues that the new Hyper algorithm for solving propositional abduction can save an exponentially larger number of iterations when compared with the AbHS/AbHS+ algorithm proposed in earlier work [58].

Consider a PAP $P_1 = (V, H, M, T, c)$, with:

$$\begin{aligned}
V &= \{t_1, x_1, y_1, m_1, \dots, t_n, x_n, y_n, m_n\} \\
H &= \{(\neg x_1), (x_1 \vee t_1), (\neg y_1), (y_1 \vee t_1), \dots, \\
&\quad (\neg x_n), (x_n \vee t_n), (\neg y_n), (y_n \vee t_n)\} \\
M &= \{(m_1), (m_2), \dots, (m_n)\} \\
T &= \{(\neg t_1 \vee \neg t_2 \vee \dots \vee \neg t_n), \\
&\quad (\neg t_1 \vee m_1), \dots, (\neg t_n \vee m_n)\}
\end{aligned} \tag{11}$$

and $c(C_i) = 1$ for $C_i \in H$. Clearly, P has no solution. For M to be entailed, S must imply all variables t_i to 1; but this causes $T \wedge S$ to become inconsistent. Moreover, there are exponentially many sets S , which are not consistent with T . AbHS+ will have to enumerate all of these sets S and, for each such set S , it will use one additional SAT oracle call to conclude that $T \wedge S$ is inconsistent. Since AbHS+ (or AbHS) selects all falsified clauses when blocking counterexamples of $T \wedge S \wedge (\neg M)$, all subsets of S inconsistent with T will be eventually enumerated. In contrast, since Hyper ensures consistency between S and T when selecting a minimum hitting set, this exponentially large number of oracle calls is not observed. (These differences between AbHS/AbHS+ and Hyper are experimentally validated in [Section 4](#).)

It should be clear that the exponentially large reduction in the number of oracle calls obtained with Hyper are hidden in the minimum hitting set extractor. However, in Hyper the minimum hitting set extractor is based on MaxSAT (concretely core-guided MaxSAT), and so this hidden complexity is handled (most often efficiently) by the SAT solver.

The inclusion of M to find each set S can also potentially save exponentially many iterations. Consider the following PAP $P_2 = (V, H, M, T, c)$:

$$\begin{aligned} V &= \{m, t_1, x_1, \dots, t_n, x_n\} \\ H &= \{(m \vee \neg x_1), (m \vee x_1 \vee t_1), \\ &\quad (m \vee \neg x_n), (m \vee x_n \vee t_n)\} \\ M &= \{(m)\} \\ T &= \{(\neg t_1 \vee \dots \vee \neg t_n)\} \end{aligned} \tag{12}$$


In contrast with the previous example, P has a solution, i.e. there exists a subset S of H (with $S = H$) such that $T \wedge S \models M$. Until the solution is found, all computed models of $T \wedge S \wedge (\neg M)$ will also falsify $T \wedge S \wedge M$. Any of these models might be filtered out if any candidate set S is such that $T \wedge S \wedge M$ is consistent. It should be noted that in this case there is no formal guarantee that the number of SAT oracle calls must be exponential. This depends on the solutions provided by the minimum hitting set algorithm used. Essentially, taking M into account when selecting S guarantees that the picked set S will not be such that $T \wedge S \models \neg M$. As the results in the next section confirm, in practice AbHS/AbHS+ can generate exponentially many candidates S for which $T \wedge S \models \neg M$.

4 Experimental Results


This section evaluates the proposed approach to propositional abduction.

4.1 Experimental Setup

All the conducted experiments were performed in Ubuntu Linux on an Intel Xeon E5-2630 2.60GHz processor with 64GByte of memory. The time limit was set to 1800s and the memory limit to 10GByte for each process to run. A prototype of the Hyper algorithm proposed above was implemented in C++ and



(a) PMS instances



(b) WPMS instances

Fig. 1: Performance of Hyper, Hyper*, AbHS, and AbHS+ for the Abduction Problem Suite benchmarks.

consists of two interacting parts. One of them computes minimum size hitting sets of the set of counterexamples, also satisfying $T \wedge S \wedge M$. This is achieved with the use of an incremental implementation of the algorithm based on soft

cardinality constraints [2, 47], which is a state-of-the-art MaxSAT algorithm that won several categories in the MaxSAT Evaluation 2015¹. The other part of the prototype checks satisfiability of $T \wedge S \wedge (\neg M)$, where S is a candidate hitting set reported by the hitting set solver. Note that both parts of the solver were implemented on top of the well-known SAT solver Glucose 3.0² [5].

Besides the basic version of Hyper, we also implemented an improved version, which is below referred to as Hyper* and contains the first two improvements described in Section 12. Namely, the first improvement does partial reduction of counterexamples by traversing and trying to satisfy $0.2 \times m$ clauses of each counterexample, where m is the size of the counterexample. The second improvement used in Hyper* consists in bootstrapping the hitting set solver with 100 MCSes of MaxSAT formula $T \wedge H \wedge (\neg M)$. It should be noted that bootstrapping the algorithm is not necessary but in some cases it can boost the performance of the main algorithm. Also note that the MaxSAT solver in both Hyper and Hyper* trims unsatisfiable cores [47] detected during the solving process at most 5 times.

The performance of Hyper and Hyper* was compared to the performance of the recent state-of-the-art algorithms AbHS and AbHS+³ [58]. Additionally, we also implemented the QMaxHS approach described in Section 3.3. The implementation was done on top of DepQBF⁴, the known QBF solver which is capable or reporting unsatisfiable cores [42]. However, the performance of QMaxHS is poor (i.e. in our evaluation it did not solve any instance from the considered benchmark suite) and we decided to exclude it from consideration.

4.2 Abduction Problem Suite

In order to assess the efficiency of the new approach to propositional abduction, the following benchmark suite was used, which was proposed and also considered in [58]. According to [58], the benchmark instances were generated based on crafted and industrial instances from MaxSAT Evaluation 2014 with the use of the MaxSAT solver LMHS⁵ [57] and the backbone solver minibones⁶ [30]. The reader is referred to [58] for details. The resulting benchmark suite contains 6 benchmarks sets: *pms-5*, *pms-10*, *pms-15*, *wpms-5*, *wpms-10*, and *wpms-15*, where the number indicates the number of manifestations. In the conducted experimental evaluation, benchmark sets were aggregated based on their type (weighted or unweighted) and resulted in two benchmark sets: PMS and WPMS, having 847 and 795 instances, respectively. The total number of instances is 1642.

The cactus plots reporting the performance of the considered algorithms measured for the considered problem instances is shown in Figure 1. As one can see in Figure 1a, for PMS benchmark instances, Hyper and Hyper* perform significantly better than both AbHS and AbHS+. More precisely, Hyper* solves 321

¹ See results for MSCG15b at <http://www.maxsat.udl.cat/15/>

² <http://www.labri.fr/perso/lsimon/glucose>

³ <http://cs.helsinki.fi/group/coreo/abhs/>

⁴ <http://lonsing.github.io/depqbf/>

⁵ <http://www.cs.helsinki.fi/group/coreo/lmhs/>

⁶ <http://sat.inesc-id.pt/~mikolas/sw/minibones/>




Fig. 2: Hyper* vs AbHS+.

Table 1: The number of iterations and running time per solver for example family (11). Additionally, for AbHS/AbHS+ the number of iterations of type 2 is also shown (in parentheses). Value n varies from 1 to 10.

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|--------------|--------|---------|-----------|------------|----------|----------|-----------|------------|--------|--------|
| AbHS | 11 (7) | 59 (49) | 363 (343) | 2401 (829) | — | — | — | — | — | — |
| | 0.0s | 0.1s | 3.4s | 166.2s | >1800s | >1800s | >1800s | >1800s | >1800s | >1800s |
| AbHS+ | 6 (2) | 14 (4) | 28 (8) | 51 (16) | 125 (32) | 388 (64) | 978 (128) | 2242 (256) | — | — |
| | 0.0s | 0.0s | 0.0s | 0.0s | 0.3s | 2.8s | 24.3s | 180.3s | >1800s | >1800s |
| Hyper | 6 | 14 | 17 | 19 | 27 | 32 | 32 | 35 | 39 | 48 |
| | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s |

instances (out of 847), which is 76 more ($> 31\%$) than the number of instances solved by AbHS+ (245). The second best competitor is Hyper, which solves 318 instances being 3 instances behind Hyper*. Finally, the worst performance is shown by AbHS, which solves 174 instances.

In contrast to the unweighted problem instances, the performance of the new algorithm for weighted instances is penalized by the use of MaxSAT for computing minimal hitting sets. The reason is that the MaxSAT solver used in Hyper does not exploit any modern and widely used heuristics to efficiently deal with weights, e.g. Boolean lexicographic optimization [43] or stratification [3].⁷

⁷ This conjecture is also suggested by the average numbers of iterations done by Hyper* and AbHS+ for the WPMS benchmarks, which are 69 and 229, respectively. Since Hyper* does significantly fewer iterations (on average) and solves around the

Table 2: The number of iterations and running time per solver for example family (12). Value n varies from 1 to 10.

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---------------|------|------|------|------|------|------|-------|--------|--------|--------|
| AbHS / | 5 | 11 | 21 | 54 | 133 | 350 | 878 | 1995 | — | — |
| AbHS+ | 0.0s | 0.1s | 0.1s | 0.2s | 0.3s | 2.0s | 15.5s | 168.8s | >1800s | >1800s |
| Hyper | 6 | 16 | 17 | 14 | 20 | 29 | 18 | 27 | 30 | 37 |
| | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s | 0.0s |

This can explain a similar performance shown by Hyper and Hyper* compared to AbHS+. More precisely, Hyper* is able to solve 398 instances (out of 795). AbHS+ comes second solving 389 instances while Hyper is 2 instances behind AbHS+ (387 solved). The worst performance is shown again by AbHS, which solves 252 instances.

Regarding the performance of the *virtual best solver* (VBS), the data for both sets of benchmarks can be seen in Figure 1 and it is the following. For PMS, the VBS aggregating Hyper, Hyper* as well as AbHS+⁸ is able to solve 328 instances, which is 7 more instances than what Hyper* can solve alone. In contrast, for the WPMS instances the picture is different: the VBS solves 425 instances, which is 27 and 36 more instances than what Hyper* and AbHS+ can solve separately. This indicates that Hyper* and AbHS+ complement each other in this case, which suggests building a portfolio of the solvers for weighted instances. The performance comparison between AbHS+ and Hyper* is detailed in the scatter plots shown in Figure 2 and also confirms this conclusion.

Although the experimental results for the abduction problem suite show clear performance gain of the proposed Hyper algorithm over the state-of-the-art in propositional abduction (AbHS+), it is important to mention that the benchmark suite was generated (see [58]) from the MaxSAT instances by filtering out those of them that are hard for MaxHS-like MaxSAT solvers, i.e. MaxHS [13, 14] and LMHS [57]. This fact suggests that applying similar ideas for generating problem instances by targeting MaxSAT formulas that are easier for another family of MaxSAT algorithms, e.g. the core-guided algorithms based on soft cardinality constraints [47] (recall that the MaxSAT solver in Hyper is one of them), would result in even better performance of Hyper. Moreover, the experimental results for the weighted benchmarks emphasize the importance of applying modern techniques (e.g. Boolean lexicographic optimization and stratification) targeting specifically weighted instances. Having implemented such improvements, one could expect Hyper to perform better for problem instances with weights.

same number of instances as AbHS+ does, we assume that the calls to the MaxSAT oracle are harder (on average).

⁸ AbHS is excluded from the VBS since it does not contribute to its performance.

4.3 Oracle Calls in AbHS+

This section studies the number of iterations for the considered approaches for the families of examples described in [Section 12](#). For both examples (11) and (12) we generated 10 instances varying size n of the instance from 1 to 10 in order to show how the number of iterations grows with the growth of size n for each approach.⁹

Recall that example (11) aims at showing the importance of adding the theory clauses into the hitting set solver by saving an exponential number of iterations related to candidates that are not consistent with the theory. In AbHS/AbHS+ the consistency check is done through the second SAT call and results in a counterexample blocking the candidate (AbHS+ also blocks its supersets). The idea of proposing example family (11) is, thus, to show that the number of iterations of this type (let us call them iterations of type 2 because they are related with the 2nd SAT call) in AbHS and AbHS+ can be exponentially larger than the number of iterations done by Hyper¹⁰. Indeed, [Table 1](#) confirms this conjecture indicating that the number of iterations of type 2 in AbHS+ grows exponentially with the growth of n , i.e. it is exactly 2^n (see the values in parentheses), while the number of iterations done by Hyper is negligible. The situation gets even more dramatic for AbHS. (As one can see, the total number of iterations performed by AbHS and AbHS+ grows even faster.) The running time spent by AbHS and AbHS+ also grows significantly with the growth of n . As a result, AbHS and AbHS+ cannot solve any instances for $n > 4$ and $n > 8$, respectively, within 1800 seconds. Observe that Hyper reports the result for each n immediately.

Regarding example (12), it shows the importance of adding M into the hitting set solver. [Table 2](#) confirms that it can save an exponential number of iterations. As one can observe, the number of iterations grows exponentially with growing value n for AbHS and AbHS+. Note that in this case AbHS and AbHS+ behave similarly to each other, which is why [Table 2](#) does not have a separate row for AbHS+. Analogously to the previous case, the performance of the solver (i.e. its running time) is severely affected by the number of iterations. Analogously to the previous example and in contrast to AbHS and AbHS+, the basic version of Hyper does significantly fewer iterations and spends almost no time for each of the considered instances.

⁹ It should be noted that the pseudo-code in [Algorithm 2](#) (taken from [58]), as well as the actual source code, needs to be modified for AbHS+ to produce correct results when a PAP does not have a solution, as is the case with example (11). When blocking previously computed hitting sets, AbHS+ can generate clauses both with positive literals and clauses with negative literals and, for a PAP without solution, it will eventually compute an empty hitting set, denoting that there is no solution to the constraints added as sets to hit. As a result, the pseudo-code (and the source) needs to test for the case when the minimum hitting set returned is empty, in which case it must return ‘no solution’. This fix was added to AbHS+ to get the results presented in this section.

¹⁰ In order to test the number of iterations, all the optimizations related to Hyper* were turned off and, thus, the basic version of Hyper was considered instead.

5 Conclusions

Abduction finds many applications in Artificial Intelligence, with a large body of work over the years. Recent work investigated propositional abduction, and proposed the use of a variant of the implicit hitting set algorithm MaxHS for solving the problem [58].

This paper identifies several sources of inefficiency with earlier work, and proposes a novel, implicit hitting set inspired, algorithm for propositional abduction. The novel algorithm, Hyper, is shown to outperform the recently proposed algorithms AbHS and AbHS+ on existing problem instances. In addition, the paper demonstrates that the proposed improvements can result in exponential savings on the number of SAT oracle calls, which helps explain the observed performance improvements. In a broader context, this paper contributes to the recent body of work on implicit hitting set algorithms, and identifies algorithmic optimizations that can be significant in other contexts.

A number of research directions can be envisioned. These include improvements to the MaxSAT solver used for computing minimum hitting sets, as this represents the main bottleneck of the Hyper algorithm. Additional work will involve applying Hyper to a larger range of problem instances.

References

1. M. Alberti, F. Chesani, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. Security protocols verification in abductive logic programming: A case study. In *ESAW*, pages 106–124, 2005.
2. B. Andres, B. Kaufmann, O. Matheis, and T. Schaub. Unsatisfiability-based optimization in clasp. In *ICLP*, pages 211–221, 2012.
3. C. Ansótegui, M. L. Bonet, and J. Levy. SAT-based MaxSAT algorithms. *Artif. Intell.*, 196:77–105, 2013.
4. M. F. Arif, C. Mencia, and J. Marques-Silva. Efficient MUS enumeration of horn formulae with applications to axiom pinpointing. In *IJCAI*, pages 324–342, 2015.
5. G. Audemard, J. Lagniez, and L. Simon. Improving Glucose for incremental SAT solving with assumptions: Application to MUS extraction. In *SAT*, pages 309–317, 2013.
6. J. Bailey and P. J. Stuckey. Discovery of minimal unsatisfiable subsets of constraints using hitting set dualization. In *PADL*, pages 174–186, 2005.
7. A. Biere, M. Heule, H. van Maaren, and T. Walsh, editors. *Handbook of Satisfiability*. IOS Press, 2009.
8. E. Birnbaum and E. L. Lozinskii. Consistent subsets of inconsistent systems: structure and behaviour. *J. Exp. Theor. Artif. Intell.*, 15(1):25–46, 2003.
9. T. Bylander, D. Allemang, M. C. Tanner, and J. R. Josephson. The computational complexity of abduction. *Artif. Intell.*, 49(1-3):25–60, 1991.
10. K. Chandrasekaran, R. M. Karp, E. Moreno-Centeno, and S. Vempala. Algorithms for implicit hitting set problems. In *SODA*, pages 614–629, 2011.
11. A. Cimatti, A. Griggio, B. J. Schaafsma, and R. Sebastiani. A modular approach to MaxSAT modulo theories. In *SAT*, pages 150–165, 2013.
12. N. Creignou and B. Zanuttini. A complete classification of the complexity of propositional abduction. *SIAM J. Comput.*, 36(1):207–229, 2006.

13. J. Davies and F. Bacchus. Solving MAXSAT by solving a sequence of simpler SAT instances. In *CP*, pages 225–239, 2011.
14. J. Davies and F. Bacchus. Postponing optimization to speed up MAXSAT solving. In *CP*, pages 247–262, 2013.
15. I. Dillig and T. Dillig. Explain: A tool for performing abductive inference. In *CAV*, pages 684–689, 2013.
16. I. Dillig, T. Dillig, and A. Aiken. Automated error diagnosis using abductive inference. In *PLDI*, pages 181–192, 2012.
17. I. Dillig, T. Dillig, B. Li, and K. L. McMillan. Inductive invariant generation via abductive inference. In *OOPSLA*, pages 443–456, 2013.
18. T. Eiter and G. Gottlob. The complexity of logic-based abduction. *J. ACM*, 42(1):3–42, 1995.
19. T. Eiter and K. Makino. Abduction and the dualization problem. In *Discovery Science*, pages 1–20, 2003.
20. B. el Ayeb, P. Marquis, and M. Rusinowitch. Preferring diagnoses by abduction. *IEEE Trans. Systems, Man, and Cybernetics*, 23(3):792–808, 1993.
21. M. Gavanelli, M. Alberti, and E. Lamma. Integrating abduction and constraint optimization in constraint handling rules. In *ECAI*, pages 903–904, 2008.
22. M. Gavanelli, M. Alberti, and E. Lamma. Integration of abductive reasoning and constraint optimization in SCIFF. In *ICLP*, pages 387–401, 2009.
23. A. Ignatiev, M. Janota, and J. Marques-Silva. Quantified maximum satisfiability: A core-guided approach. In *SAT*, pages 250–266, 2013.
24. A. Ignatiev, M. Janota, and J. Marques-Silva. Quantified maximum satisfiability. *Constraints*, 21(2):277–302, 2016.
25. K. Inoue. Automated abduction. In *Computational Logic: Logic Programming and Beyond, Essays in Honour of Robert A. Kowalski, Part II*, pages 311–341, 2002.
26. K. Inoue, A. Doncescu, and H. Nabeshima. Completing causal networks by meta-level abduction. *Machine Learning*, 91(2):239–277, 2013.
27. K. Inoue, T. Sato, M. Ishihata, Y. Kameya, and H. Nabeshima. Evaluating abductive hypotheses using an EM algorithm on BDDs. In *IJCAI*, pages 810–815, 2009.
28. M. Janota, W. Klieber, J. Marques-Silva, and E. M. Clarke. Solving QBF with counterexample guided refinement. In *SAT*, pages 114–128, 2012.
29. M. Janota, W. Klieber, J. Marques-Silva, and E. M. Clarke. Solving QBF with counterexample guided refinement. *Artif. Intell.*, 234:1–25, 2016.
30. M. Janota, I. Lynce, and J. Marques-Silva. Algorithms for computing backbones of propositional formulae. *AI Commun.*, 28(2):161–177, 2015.
31. M. Janota and J. Marques-Silva. Abstraction-based algorithm for 2QBF. In *SAT*, pages 230–244, 2011.
32. M. Janota and J. Marques-Silva. Solving QBF by clause selection. In *IJCAI*, pages 325–331, 2015.
33. E. S. Jr. A linear constraint satisfaction approach to cost-based abduction. *Artif. Intell.*, 65(1):1–28, 1994.
34. A. C. Kakas, R. A. Kowalski, and F. Toni. Abductive logic programming. *J. Log. Comput.*, 2(6):719–770, 1992.
35. A. C. Kakas and P. Mancarella. Database updates through abduction. In *VLDB*, pages 650–661, 1990.
36. A. C. Kakas and F. Riguzzi. Learning with abduction. In *ILP*, pages 181–188, 1997.
37. R. M. Karp. Implicit hitting set problems and multi-genome alignment. In *CPM*, page 151, 2010.

38. M. H. Liffiton and A. Malik. Enumerating infeasibility: Finding multiple MUSes quickly. In *CPAIOR*, pages 160–175, 2013.
39. M. H. Liffiton, A. Previti, A. Malik, and J. Marques-Silva. Fast, flexible MUS enumeration. *Constraints*, 21(2):223–250, 2016.
40. M. H. Liffiton and K. A. Sakallah. Algorithms for computing minimal unsatisfiable subsets of constraints. *J. Autom. Reasoning*, 40(1):1–33, 2008.
41. F. Lin and J. You. Abduction in logic programming: A new definition and an abductive procedure based on rewriting. *Artif. Intell.*, 140(1/2):175–205, 2002.
42. F. Lonsing and U. Egly. Incrementally computing minimal unsatisfiable cores of QBFs via a clause group solver API. In *SAT*, pages 191–198, 2015.
43. J. Marques-Silva, J. Argelich, A. S. Graca, and I. Lynce. Boolean lexicographic optimization: algorithms & applications. *Annals of Mathematics and Artificial Intelligence (AMAI)*, 62(3-4):317–343, 2011.
44. J. Marques-Silva, F. Heras, M. Janota, A. Previti, and A. Belov. On computing minimal correction subsets. In *IJCAI*, 2013.
45. E. Moreno-Centeno and R. M. Karp. The implicit hitting set approach to solve combinatorial optimization problems with an application to multigenome alignment. *Operations Research*, 61(2):453–468, 2013.
46. A. Morgado, F. Heras, M. H. Liffiton, J. Planes, and J. Marques-Silva. Iterative and core-guided maxsat solving: A survey and assessment. *Constraints*, 18(4):478–534, 2013.
47. A. Morgado, A. Ignatiev, and J. Marques-Silva. MSCG: Robust core-guided MaxSAT solving. *JSAT*, 9:129–134, 2015.
48. A. Morgado, M. H. Liffiton, and J. Marques-Silva. MaxSAT-based MCS enumeration. In *HVC*, pages 86–101, 2012.
49. G. Nordh and B. Zanuttini. What makes propositional abduction tractable. *Artif. Intell.*, 172(10):1245–1284, 2008.
50. A. Pfandler, R. Pichler, and S. Woltran. The complexity of handling minimal solutions in logic-based abduction. *J. Log. Comput.*, 25(3):805–825, 2015.
51. R. Pichler and S. Woltran. The complexity of handling minimal solutions in logic-based abduction. In *ECAI*, pages 895–900, 2010.
52. D. Poole. Probabilistic horn abduction and bayesian networks. *Artif. Intell.*, 64(1):81–129, 1993.
53. A. Previti and J. Marques-Silva. Partial MUS enumeration. In *AAAI*, 2013.
54. O. Ray and K. Inoue. A consequence finding approach for full clausal abduction. In *Discovery Science*, pages 173–184, 2007.
55. J. A. Reggia, B. T. Perricone, D. S. Nau, and Y. Peng. Answer justification in diagnostic expert systems-Part I: Abductive inference and its justification. *IEEE Trans. Biomedical Engineering*, (4):263–267, 1985.
56. R. Reiter. A theory of diagnosis from first principles. *Artif. Intell.*, 32(1):57–95, 1987.
57. P. Saikko, J. Berg, and M. Järvisalo. LMHS: A SAT-IP hybrid MaxSAT solver. In *SAT*, 2016.
58. P. Saikko, J. P. Wallner, and M. Järvisalo. Implicit hitting set algorithms for reasoning beyond NP. In *KR*, 2016.
59. C. Sakama and K. Inoue. An abductive framework for computing knowledge base updates. *TPLP*, 3(6):671–713, 2003.
60. K. Satoh and T. Uno. Enumerating minimal explanations by minimal hitting set computation. In *KSEM*, pages 354–365, 2006.
61. R. T. Stern, M. Kalech, A. Feldman, and G. M. Provan. Exploring the duality in conflict-directed model-based diagnosis. In *AAAI*, 2012.